

Abstract

By transforming a query into a product of conditional selectivity expressions, an existing set of statistics on query expressions can be used more effectively to estimate cardinality values. Conditional selectivity values are progressively separated according to rules of conditional probability to yield a set of non-separable decompositions that can be matched with the stored statistics on query expressions. The stored statistics are used to estimate the selectivity of the query and the estimated selectivity can be multiplied by the Cartesian product of referenced tables to yield a cardinality value.